idr
Internet-Draft
Intended status: Standards Track
Expires: April 12, 2020

J. Hu
Nokia

October 10, 2019

BGP Provisioned IPsec Transport Mode Protected Tunnel Configuration
draft-hujun-idr-bgp-ipsec-transport-mode-00

Abstract

This document defines a method of using BGP to advertise IPsec transport mode protected tunnel (like GRE tunnel with IPsec transport mode protection) configuration along with NLRI, based on [I-D.ietf-idr-tunnel-encaps] and [I-D.hujun-idr-bgp-ipsec].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 12, 2020.

Table of Contents

1. Introduction 2
1.1. Terminology 2
1.2. IPsec Transport Protected sub-TLV 3
2. Semantics and Operation 3
3. IANA Considerations 5
4. Security Considerations 5
5. Change Log 6
6. References 6
6.1. Normative References 6
6.2. Informative References 6
Author's Address 7

1. Introduction

[I-D.hujun-idr-bgp-ipsec] defines a method of using BGP to advertise configuration for IPsec tunnel with ESP tunnel mode, however there are other use cases require of using IPsec/ESP transport mode with other types of IP tunnel, like GRE tunnel, as defined in [RFC4301] and [RFC4303]. Figure 2 shows an example of IPv4 GRE tunnel packet with ESP transport mode protection. This document defines a method of using BGP to advertise configuration for these use cases.

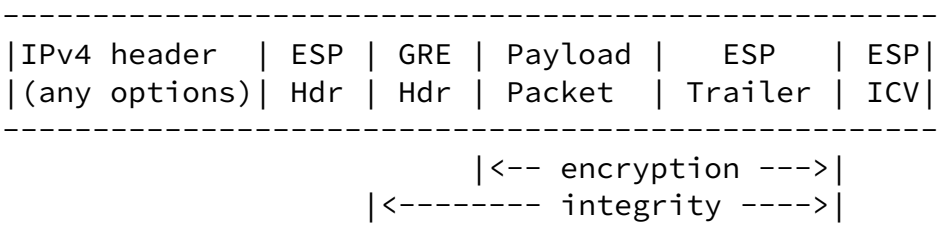

Figure 1: IPv4 GRE tunnel packet with ESP transport protection

The method follows same principle as [I-D.hujun-idr-bgp-ipsec], keep changes to BGP minimal and not changing IKEv2/IPsec; however the IPsec transport mode protected IP tunnel is not a tunnel stack or nested tunnels, IPsec transport mode protection doesn't add extra IP header.

The requirement of using IPsec transport mode is signaled by including a sub-TLV: IPsec transport protected, in a BGP tunnel encapsulation TLV.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP

14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. IPsec Transport Protected sub-TLV

This sub-TLV represents using IPsec transport mode protection for the tunnel specified by parent tunnel encapsulation TLV, its value is a IPsec configuration tag as defined in [I-D.hujun-idr-bgp-ipsec].

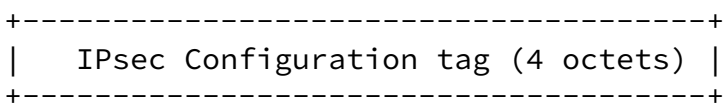

Figure 2: IPsec Configuration Tag

For a given tunnel encapsulation TLV, IPsec configuration tag sub-TLV MUST appear only one time.

2. Semantics and Operation

Except for what this document explicitly specifies, the semantics and operation of tunnel encapsulation TLV with IPsec Transport Protected sub-TLV are same as defined in [I-D.ietf-idr-tunnel-encaps] and [I-D.hujun-idr-bgp-ipsec].

IPsec Transport Protected sub-TLV MAY be included in any type of IP tunnel TLV specified in [I-D.ietf-idr-tunnel-encaps]; it MUST be ignored when included in a IPsec tunnel TLV.

The inclusion of IPsec Transport Protected TLV and its value is determined by local policy.

Following are the rules of operations:

1. All routers are pre-provisioned with Mapping between IPsec configuration tag value and IPsec configurations include

authentication method/credentials

2. If a given NLRI needs a specific tunnel encapsulation with IPsec transport mode protection, then advertising router need to include an IPsec Transport Protected sub-TLV with required configuration tag, in the corresponding tunnel encapsulation TLV/ attribute, along with the NLRI in BGP UPDATE U;

3. When a router need to forward a packet along a path is determined by a BGP UPDATE which has a tunnel encapsulation attribute that contains one or more tunnel TLV, router selects a tunnel TLV based on Semantics defined in [I-D.ietf-idr-tunnel-encaps], if

the selected tunnel TLV contains IPsec Transport Protected sub-TLV, then the router use first feasible CHILD_SA for IP tunnel packet encryption, a CHILD SA is considered as feasible when it meets all following conditions:

* it is ESP transport mode

* its private and public routing instance is same as routing instance in which the packet to be forwarded

* its peer tunnel address is same as indicated by Remote Endpoint sub-TLV

* the source and destination address of the packet to be forwarded falls in the range of CHILD SA's traffic selector

* its transform and other configuration maps to the tag indicated in the IPsec configuration tag sub-TLV

4. If router can't find such CHILD SA, then it will use IKEv2 to create one with following IPsec configuration:

* ESP transport mode

* private and public routing instance is the routing instance in which the packet to be forwarded

* peer tunnel address is specified by Remote Endpoint sub-TLV

    * local traffic selector:

      + address range: local tunnel endpoint address

      + protocol: tag mapped configuration

      + port range: tag mapped configuration

    * remote traffic selector:

      + address range: address in Remote Endpoint sub-TLV of selected tunnel encapsulation TLV

      + protocol: tag mapped configuration

      + port range: tag mapped configuration

    * other configurations come from mapping of the configuration tag in IPsec Transport Protected sub-TLV of selected tunnel encapsulation TLV

3. IANA Considerations

This document will request new values in IANA "BGP Tunnel Encapsulation Attribute Sub-TLVs" registry for IPsec Transport Protected sub-TLV.

4. Security Considerations

IKEv2 is used to create IPsec tunnel, which ensures following:

- Traffic protection keys are generated dynamically during IKEv2 negotiation, only known by participating peer of the IPsec tunnel; there is no central node to manage and distribute all keys.

- IKEv2 rekey mechanism refresh keys regularly; PFS(Perfect Forward Secrecy) provides additional protection;

- Secure authentication mechanism that only allow authenticated peer

to create tunnel

- Traffic Selector guarantee that only agreed traffic is allowed to be forwarded within the IPsec tunnel;

- Using a separate, dedicate protocol(IKEv2) for key management/ authentication ensure they are not tied to BGP, all existing and future IKEv2 features could be used without changing BGP;

There is concern that malicious party might manipulate IPsec tunnel encapsulation attribute to divert traffic, however this risk could be mitigated by IKEv2 mutual authentication.

BGP route filter include outbound route filter [RFC5291], Origin Validation [RFC6811] and BGPSec [RFC8205] could be used to further secure BGP UPDATE message.

IKEv2 cookie [RFC7296] and varies mechanisms defined including client puzzle defined in [RFC8019] could be used to protect IKEv2 from Distributed Denial-of-Service Attacks.

Follow latest IETF ESP/IKEv2 implementation requirement and guidance ([RFC8221] and [RFC8247] at time of writing) to make sure always using secure and up-to-date cryptographic algorithms;

5. Change Log

- v00 Sep 29, 2019: initial draft

6. References

6.1. Normative References

[I-D.hujun-idr-bgp-ipsec]
Hu, J., "BGP Provisioned IPsec Tunnel Configuration", draft-hujun-idr-bgp-ipsec-01 (work in progress), September 2019.

[I-D.ietf-idr-tunnel-encaps]
Patel, K., Velde, G., and S. Ramachandra, "The BGP Tunnel

Encapsulation Attribute", draft-ietf-idr-tunnel-encaps-14 (work in progress), September 2019.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <https://www.rfc-editor.org/info/rfc4301>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <https://www.rfc-editor.org/info/rfc4303>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

6.2. Informative References

[RFC5291] Chen, E. and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", RFC 5291, DOI 10.17487/RFC5291, August 2008, <https://www.rfc-editor.org/info/rfc5291>.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <https://www.rfc-editor.org/info/rfc6811>.

[RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <https://www.rfc-editor.org/info/rfc7296>.

[RFC8019] Nir, Y. and V. Smyslov, "Protecting Internet Key Exchange Protocol Version 2 (IKEv2) Implementations from Distributed Denial-of-Service Attacks", RFC 8019,

DOI 10.17487/RFC8019, November 2016,
<https://www.rfc-editor.org/info/rfc8019>.

[RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <https://www.rfc-editor.org/info/rfc8205>.

[RFC8221] Wouters, P., Migault, D., Mattsson, J., Nir, Y., and T. Kivinen, "Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 8221, DOI 10.17487/RFC8221, October 2017, <https://www.rfc-editor.org/info/rfc8221>.

[RFC8247] Nir, Y., Kivinen, T., Wouters, P., and D. Migault, "Algorithm Implementation Requirements and Usage Guidance for the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 8247, DOI 10.17487/RFC8247, September 2017, <https://www.rfc-editor.org/info/rfc8247>.

Author's Address

Hu Jun
Nokia
777 East Middlefield Road
Mountain View CA 95148
United States

Email: jun.hu@nokia.com